The SAL Architecture: A Synthesis of ACT-R and Leabra
DARPA BICA Phase I Final Report

Prepared by the ACT-R & Leabra Teams

Introduction

The goal of the BICA program is to produce flexible systems that capture the power of human cognition – systems that can both adapt to new environments and be tasked with new instructions without reprogramming. To achieve that goal we are bringing together ACT-R, a high functionality architecture, that already performs well on the goal of taskability, with Leabra, a high neural-fidelity architecture, that already performs well on adapting to new environments. The following are examples of high-end applications of ACT-R:

1. A system that can take instruction for a new domain of mathematics and reproduce student behavior (including brain imaging data) with no special programming.
2. A MOUT (Military Operations in Urban Terrain) system that simulates soldiers navigating in urban terrain and executing USMC combat doctrine while remaining adaptive to environment and opponent actions.
3. A driving model that can drive a car (in a simulator) and predict the degradation on performance that occurs when devices like cell phones or GPS systems are introduced.

The following are examples of high-end applications of Leabra:

1. A model of the visual neural pathways supporting object recognition, that learns to recognize objects despite wide variations in size, location, orientation, and context, performing at state-of-the-art levels based solely on existing general-purpose Leabra learning mechanisms.
2. A model of the prefrontal cortex and basal ganglia that can learn a wide range of complex working memory and cognitive control tasks based on trial and error learning, without any task-specific pre-configuration.
3. A model of the hippocampus and neocortex that can account for learning and memory data from over 20 different experiments on rats and humans, accurately capturing the effects of hippocampal damage and the complex division of labor between these neural systems in learning.

Although they are at different levels of description, ACT-R and Leabra have deep compatibilities that enable them to be synthesized into a new system that we will call SAL (Synthesis of ACT-R and Leabra). We have already put components of the two architectures together, and they interact successfully to solve a few interesting problems. In our work in Phase II we hope to go beyond just piecing together the best of both systems. We will combine the insights of each system into the components of the new SAL architecture. We will also develop more powerful means of interaction among the components.

In this document, we begin with a broad overview of the principles and neural basis of this new SAL architecture and how it relates to the existing ACT-R and Leabra architectures. Then, we describe a concrete instantiation of a SAL model in a simplified version of the "Egg Hunt" task, operating in the Unreal Tournament simulation engine. This model shows how SAL can do things that neither Leabra nor ACT-R can currently do by itself, demonstrating the promise of the integrated architecture. This example also shows how we can easily extend the functionality of SAL without reprogramming the system. Finally, we outline a trajectory for further development of the SAL architecture as we move into Phase II, highlighting some of the key scientific and technical challenges and payoffs. It should already be quite evident that this synthesis represents an important new development in the field of cognitive architectures, and we are only at the very early stages.

Overview of SAL

The ACT-R and Leabra architectures are both characterized by the attempt to account for a wide range of cognitive and neural phenomena using a small and therefore strongly constrained set of computational primitives, as contrasted with the predominant "one-off" and "grab-bag" cognitive models in the field. These architectures have been focused on largely complementary domains: Leabra on the neural mechanisms subserving the processing of individual stimuli and short sequences thereof, and ACT-R on more abstract, longer time-scale controlled cognition unfolding over minutes. These architectures are each arguably the most successful in their domain at rigorously accounting for a wide range of cognitive and neural phenomena, with each model providing detailed accounts of hundreds of distinct types of data.

Given their independent success at describing the human cognitive system, it is reassuring, and quite remarkable, that they have arrived at very convergent views of the overall cognitive architecture. This convergence is particularly significant given that the Leabra architecture is derived from more bottom-up neuro-computational constraints about the kinds of processing different parameterizations of a common neural substrate can support, while ACT-R is derived more top-down based on regularities and constraints present in human cognitive performance. This independent convergence provides a strong basis for confidence in the veracity of the emerging SAL architecture.

Both the Leabra and ACT-R architectures can be described at the most abstract level in terms of complimentary systems that are specialized for cognitively and neurally dissociable forms of processing. These dissociable neural systems form the basis of the SAL architecture, and can be categorized most broadly in a tripartite architecture, as previously documented by the Leabra team (see Figure 1):

- The **posterior cortex**, which performs basic sensory (e.g., visual, auditory, and somatosensory) processing (in the occipital and inferiotemporal lobes) and motor processing (in the parietal lobe, which interacts strongly with posterior frontal cortical motor areas). This area is also critical for encoding higher-level semantic and declarative knowledge about the world, including many aspects of language and reasoning (in higher level association cortex in both temporal and parietal lobes).
- The **prefrontal cortex**, which is necessary for active maintenance of information and executive control of cognitive processing, and interacts closely with the **basal ganglia,** which is specialized for action selection and learning about which actions lead to reward or punishment. This system is critical for procedural processing and learning.
- The **hippocampus**, which is responsible for rapid learning of new information, often of a

declarative (verbally-mediated) form (e.g., the location of a given object in an environment, the name of someone you've just met, or a new fact such as "the capital of Pakistan is Karachi").

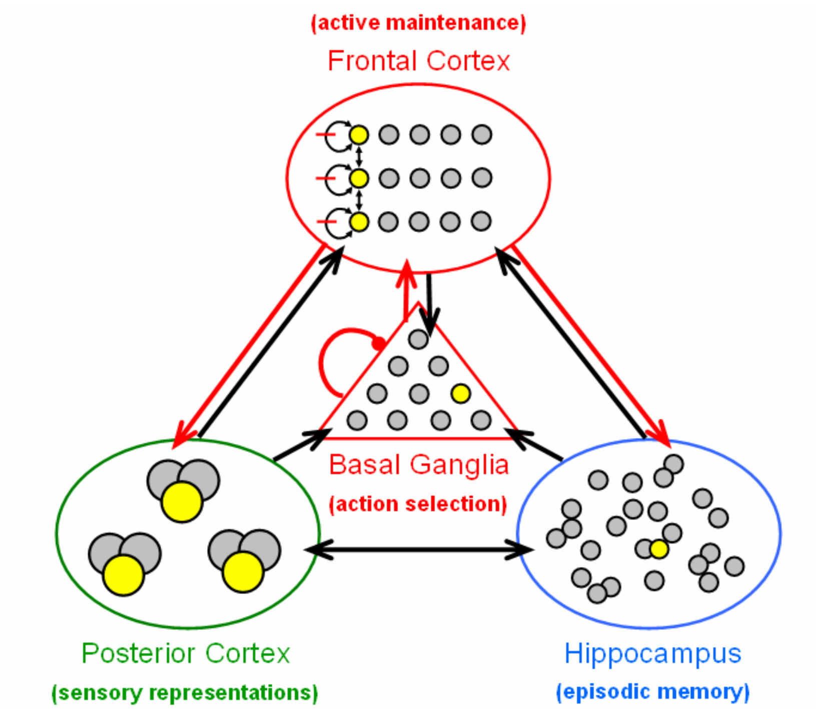

Figure 1: Tripartite Architecture of SAL. Human cognition is conceptualized in terms of the computational properties of distinct brain areas, each specialized for different incompatible forms of learning (e.g., rapid learning in the hippocampus vs. slow learning in the cortex). Red arrows represent top-down cognitive control (which results from interactions between Frontal Cortex and Basal Ganglia), while black arrows represent standard neural communication.

This tripartite, neurally-focused architecture can be decomposed into separable cognitive modules, which corresponds very closely with the ACT-R architecture, as shown in Figure 2. This mapping of functional modules onto neural structures is only approximate, particularly in the case of the imaginal and declarative modules. In the case of the imaginal module, while the control and maintenance are believed to be in the prefrontal cortex, the actual imaginal transformations seem to be performed in the parietal cortex, which is part of posterior cortex. In the case of declarative memory, while the hippocampus is a critical component, as in Leabra, much of the cortex can store declarative memory as well, and the prefrontal cortex plays an important role in controlling encoding and retrieval operations.

Before elaborating the cognitive and neural synergies between the Leabra and ACT-R architectures as captured in SAL, we can illustrate the general operation of this system in the context of SAL performing the "Egg Hunt" task in the Unreal Tournament (UT) environment (more details are provided below). SAL first hears a command like "find the armor", which initiates a search through a set of rooms until it finds the target object, at which point it takes possession of the object. In this scenario, the different modules play the following roles:

1. The **aural module** holds a representation of the spoken sentence for processing.
2. The **goal module** maintains a representation of the activity (“find”) and the object (“armor”) throughout the episode.
3. The **declarative module** is accessed both to retrieve knowledge of the layout of the rooms and to maintain knowledge of which rooms have been searched.
4. The **imaginal module** is used to maintain a representation of the current room and the locations that have already been examined in the room.
5. A **motor module** is used to request movements from room to room and to orient to various objects in the room.
6. A **visual module** is used to represent the visual scene and identify objects.
7. The **procedural module** steps the agent through the tasks of planning its moves, performing the actions, and recognizing when the task objective has been achieved. It learns to improve its performance in future attempts, based on the success and failure of these actions.

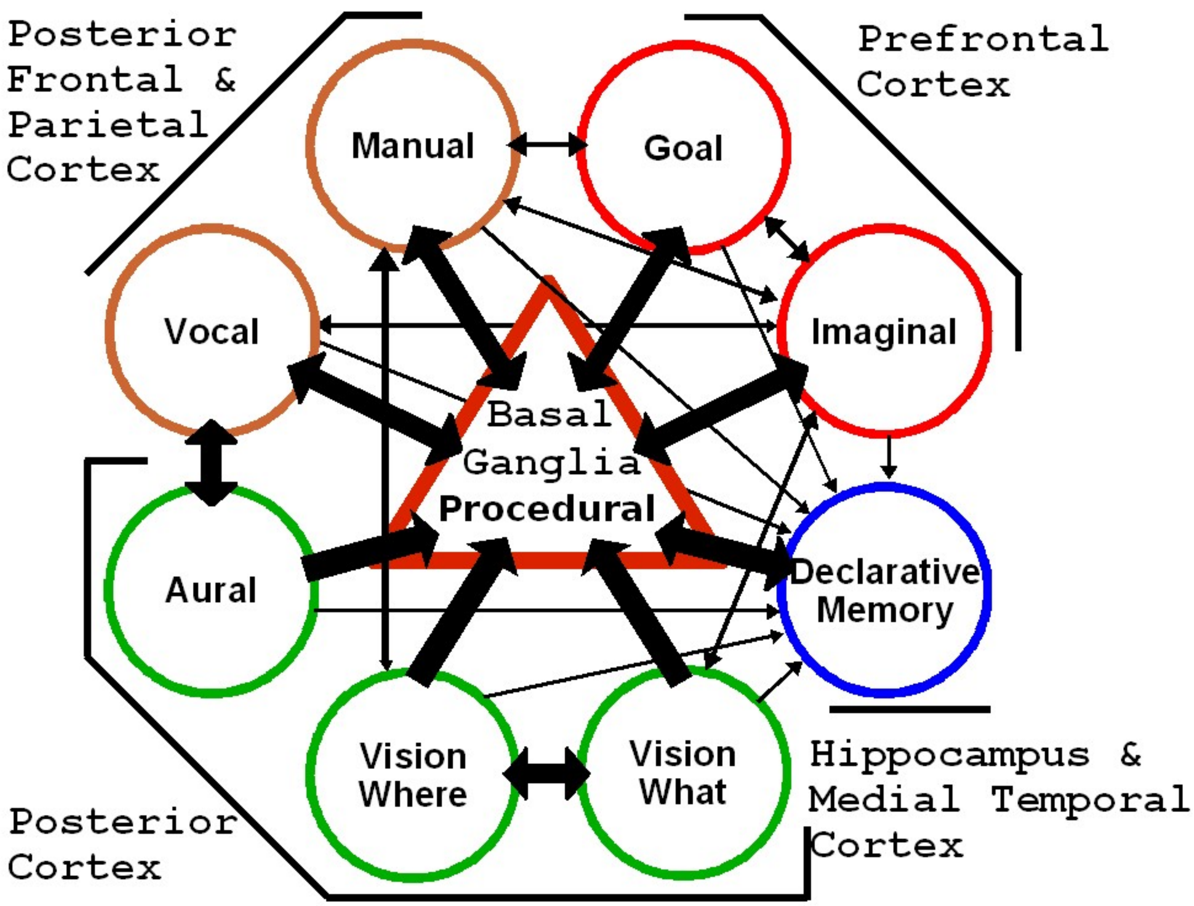

Figure 2: Cognitive Module Architecture of SAL, where the broad tripartite architecture has been subdivided into finer grained separable cognitive mechanisms.

Cognitive and Neural Synergies between ACT-R and Leabra in the SAL

Architecture

The Procedural-Declarative Distinction

The most notable area of convergence between ACT-R and Leabra is in the broad division of the cognitive architecture into procedural and declarative components. From this one distinction, many others follow, as elaborated in subsequent subsections. This distinction has clear cognitive and neural validity. People can possess abstract declarative knowledge of how to do something but be procedurally incapable of doing so (e.g., new drivers or golf players), and vice-versa (e.g., touch typists often cannot recall where the keys are located). Neurally, the basal ganglia are critical for initiating procedural actions, whereas the cortex and hippocampus support declarative knowledge.

In the Leabra framework, different types of processing are supported by the neural specializations present in the basal ganglia, compared with those present in the hippocampus and cortex. The basal ganglia system is strongly modulated by dopamine, which signals reward and punishment information. Positive reward reinforces associated procedural actions, while negative feedback reduces the likelihood of producing associated actions. A similar, more abstract form of reinforcement learning is present in the ACT-R procedural system.

On the other hand, the neural properties of the hippocampus have been shown in the Leabra framework to be critical for the rapid learning of new arbitrary information without interfering with existing knowledge. Specifically, having a relatively few neurons active at one time ("sparse representations") causes neural representations to separate from each other, minimizing interference. This rapidly acquired knowledge can, over time, be integrated into more overlapping, distributed representations in cortical areas, supporting the ability to draw sophisticated inferences and generalize to novel situations. The declarative system in ACT-R integrates both of these properties: new chunks of knowledge, encoded as combinations of existing chunks, can be rapidly formed; chunks that are used more frequently over time gain higher levels of activation and correspond to more expert knowledge; similarities can be defined between symbolic chunks to drive semantic generalization to related situations.

Although dissociable, the procedural and declarative systems interact intimately in any complete cognitive process. In ACT-R, the firing of productions is driven by the active contents of the declarative and other memory buffers, and the result of production firing is the updating of these buffers. In Leabra, the basal ganglia procedural system is tightly linked with the prefrontal cortex, which maintains task-relevant information in an active state over time. One of the primary functions of the basal ganglia in the brain is to drive the updating of these prefrontal active memory states. These prefrontal areas then influence activation states throughout the rest of the cortex via strong top-down excitatory projections. Each area of posterior cortex has an associated prefrontal area, with which it has strong bidirectional excitatory connectivity. Thus, we associate the buffers of ACT-R with these prefrontal representations of corresponding posterior cortical areas.

Reinforcement Learning for Procedural Processing

Both ACT-R and Leabra include reinforcement learning mechanisms to shape the procedural processing system. This form of learning uses success and failure information to shape the probability of selecting a given action in the future, and is dissociable from the form of learning that shapes cortical and hippocampal declarative representations. Although the detailed equations differ, there is considerable similarity between the two architectures in the computational principles underlying this learning, and both agree that the basal ganglia are its central neural locus.

Declarative Learning and Processing Mechanisms

Both Leabra and ACT-R make use of Hebbian-style learning mechanisms to modulate the strength of representations in declarative memory. Such learning mechanisms are based on the history of activation of the information stored in declarative memory; but critically, not on the success or failure of a particular action taken using that memory. This fact clearly dissociates these mechanisms from procedural reinforcement learning, and numerous cognitive experiments have validated this property of declarative memory.

In terms of processing information already stored in declarative memory, the concept of spreading activation is critical to both architectures. In ACT-R, activation spreads among declarative chunks in proportion to their associative strength. In Leabra, a similar activation spreading dynamic occurs, in that coarse-coded distributed representations in posterior cortical areas cause associated representations to overlap and share activation states.

Visual Mechanisms

With respect to vision we can distinguish between visual perception and visual attention. On the visual perception front Leabra offers a detailed theory and will be important for parsing the raster format that is anticipated for BICA Phase II. Efforts to incorporate direct perception into ACT-R have been limited to date, although there have been some proposals for extensions with systems like Robert St. Amant's "Segman." We will explore possible synergies in these approaches in the future.

ACT-R has a moderately functional overall theory of top-down attention that has been applied to vision, and which we will explore in connection with Leabra attentional processes. For example, Mike Byrne has been developing a "rational analysis of attention" on the premise that the system attempts to maximize the information uptake. This allows background biases such as a preference for rare objects to be combined with an immediate and explicit desire to find a red object. We will research how Byrne's equations for the salience of an object (a lot like activation equations in declarative memory) map to Leabra. Separately, Dario Salvucci has developed a rather sophisticated theory of eye movements for ACT-R. Salvucci's equations relate information uptake and probability of a movement to foveal distance, and also deal with the timing of saccadic programming. As Leabra visual processing is strongly dependent on successful foveation, a mapping of this theory onto Leabra control mechanisms could be quite powerful.

Initial Concrete Implementation of SAL

The SAL team has built a demonstration model representing a preliminary synthesis of the two architectures. Our goal was to anticipate the challenges we will face in implementing a truly integrated and embodied architecture for Phase II. This demonstration performs a simple version of the “Easter Egg Hunt” challenge suggested by discussions of the Phase II test problems; specifically, the SAL agent searches for a target object within a familiar environment. For this demonstration, we adopted a “first order” form of architectural integration, whereby one of the cognitive modules in an ACT-R model is replaced with a Leabra network. In this case, the Leabra model is capable of processing raw bitmap images in a way that the ACT-R model was not capable of doing; similarly, extant Leabra models are not capable of organizing problem solving behavior over a period of several minutes, as required to search for the target object in a complex environment. Thus, this SAL model represents a new level of functionality that goes beyond the capabilities of its constituent architectures. Given that this is the simplest form of integration, we are optimistic that much more interesting and powerful forms of cognition can be captured as our integration efforts develop further.

It is also worth noting that very little new work was required to make this model operational. We had already established a mechanism for ACT-R and Leabra interactions in preparation for demonstrations at the August technical meetings, including an attentional blink model and a model of the Haimson radar search task. In the attentional blink model, we combined the top-down control capabilities of ACT-R with the graded visual representations of Leabra, allowing us to account for aspects of psychological phenomena that neither architecture demonstrates individually.. In the Haimson radar search task, ACT-R and Leabra shared a symbolic representation (a name) for objects of interest and interacted dynamically. Using an existing ACT-R model for searching environments and the Leabra model of visual object recognition, we simply adapted the ACT-R task instructions and trained the Leabra model on relevant visual stimuli.

In the demonstration, the SAL agent is embodied within an Unreal Tournament simulation environment consisting of three rooms containing three categories of objects. It is familiar with the environment in that it has access to navigation points and object location points in symbolic form. Further, it has been trained to perceptually identify the three object categories from a variety of viewing angles and distances. An operator instructs SAL to find the desired target via a typed command (“find armor”). SAL must then navigate the rooms; view and perceptually identify each object; and when it recognizes the desired target, navigate to it, and picks it up.

As noted, this combined model is implemented using the Leabra system as a perceptual front-end for ACT-R, by effectively replacing the ACT-R “Vision What” module; viewed conversely, ACT-R serves as a top-down control for the Leabra vision model. ACT-R must decide which navigation point to visit next and which objects to view from that navigation point. Upon selection of the object to view, ACT-R provides the digital image that is a snapshot of the view of that object from SAL’s current location. Leabra attempts to identify the object in the center of the image, and responds with its conclusion in symbolic form. If the object matches the specified target, ACT-R navigates to it, picks it up, and navigates back to the starting point; otherwise it selects a new object to view or a new location to which to navigate.

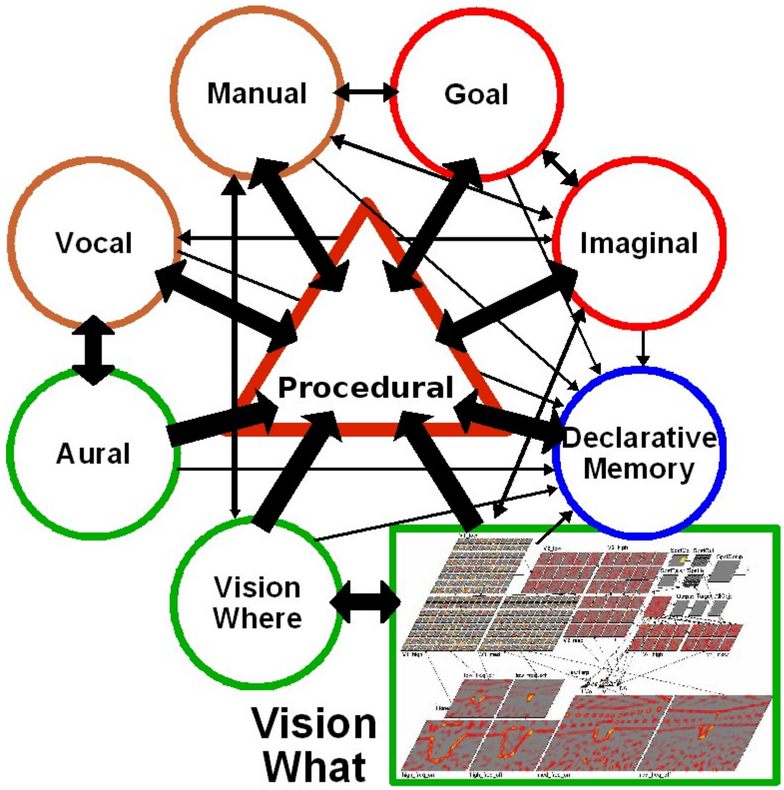

Figure 3: The SAL demonstration model architecture. Navigation through the known environment is performed by the ACT-R component, and visual object identification is handled by a Leabra neural network component. The neural network replaces ACT-R's "vision: what" module; the ACT-R production system replaces Leabra's prefrontal cortex / basal ganglia element.

The ACT-R component of the model deals with the challenges of performing the search in an efficient manner. It plans where to begin searching, and uses its episodic memory and inhibition of return capabilities to remember where it has already searched. The Leabra component of the model was trained to recognize the object categories by repeated presentations of each object from a variety of perspectives and distances and in different room backgrounds. Its non-embodied performance on novel examples is 96%, when using a variety of backgrounds and object angles; simplifying the object angles and ensuring perfect foveation increases its performance to 100%.

This model provides a number of benefits:

- It illustrates a simple connection of the two architectures operating together, and demonstrates that it is possible to bridge the gap between their different levels of description.
- It serves as a first "embodiment" of a combined model operating in a simulation environment.
- It demonstrates that the system can adapt to new environments and be tasked with new instructions without reprogramming.

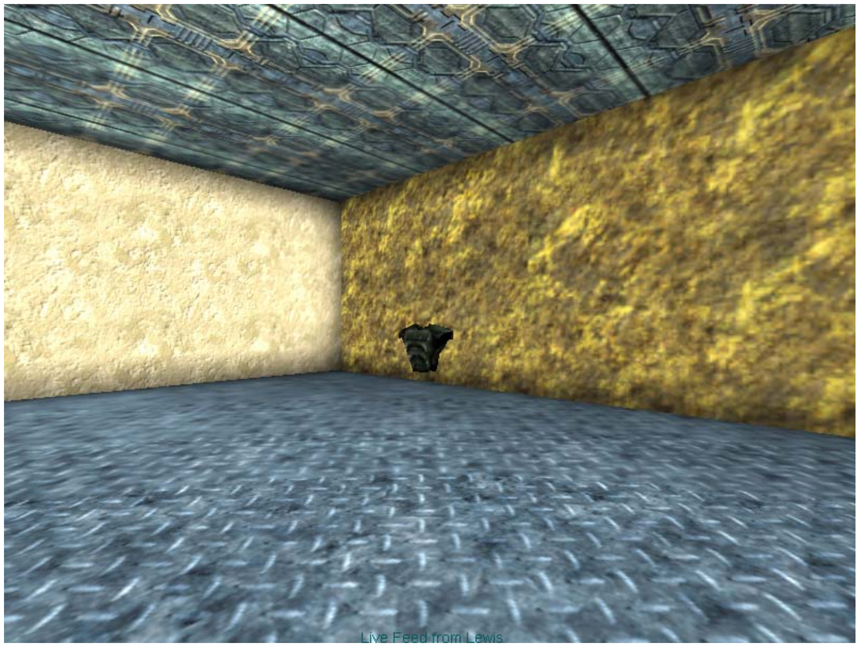

Figure 4: A sample view of a room with single object in the virtual environment.

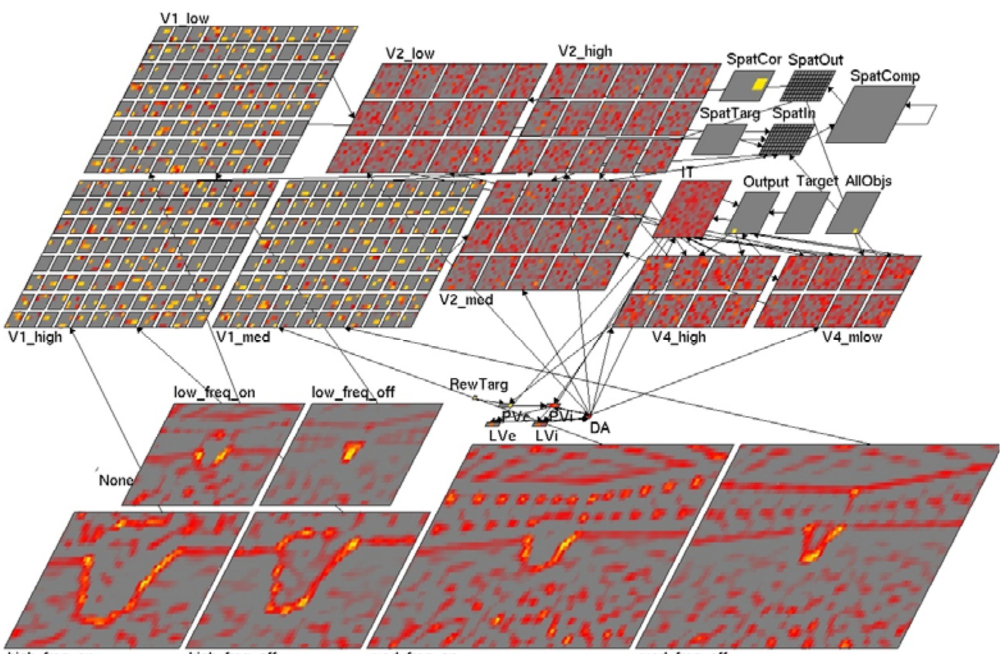

Figure 5: The state of the Leabra network as it finishes correctly identifying the armor in the scene above. The network mimics the properties of the visual system, such as: a hierarchy of visual areas, with shortcut connections between them; representions of the image at several scales; and learning from visual experience.

By combining the two architectures in a straightforward way in the context of this task, the immediate future research efforts are obvious and compelling. There are a number of clear directions for improvement:

1. The combined model must use its “Visual What” system for navigation as well as object identification. It must learn to recognize doorways, room corners, obstacles, and other important navigational cues, and the control system must learn what to do with these inputs.
2. Rather than directly specifying a next location by its coordinates, the system must navigate using simpler locomotive operations, such as “move forward” or “turn right.” These operations could then be further elaborated to interact with more realistic effectors by a motor module based on previous Leabra work on learning motor movements. This was done in the ACT-R MOUT system, and that model can be adapted for SAL.
3. Rather than automatically orienting directly toward an object or navigational cue, the combined model must perform visual search in a way that is guided by both perceptual inputs and task demands. This will require a deeper and more sophisticated integration of ACT-R and Leabra, combining body movements and saccades with dual goals of learning the navigation options and attempting to find and identify objects in the environment. This deeper integration is likely to involve integrating ACT-R models of visual salience and eye movements with Leabra calculations graded activation information.
4. The system must be able to purposefully learn about new perceptual inputs. The Leabra perceptual system will need to recognize novelty in objects, actions, or outcomes, and the cognitive system will respond by making a goal of learning more about the novel element and taking active steps to do so. This approach is inspired by the theory that human infants act as scientists, making hypotheses and performing simple experiments to test them.
5. Since navigation points and object locations will no longer be available in symbolic form, the model’s internal representations of “where it has been” and “where it has looked” must instead be connected to perceptual representations of elements or items identified in the environment. This will also drive a deeper integration of ACT-R and Leabra, by requiring either a neurally-based episodic memory, symbolic chunks that have non-symbolic perceptual components, or both.
6. The control model must become more robust to perceptual and motor errors. Since the visual system will sometimes identify or search incorrectly, and the motor system may bump into obstacles or walls, the system must recover from these errors and choose a viable strategy for recovering and continuing. The ACT-R community has had some experience with achieving this robustness in the MOUT system and we believe this experience can be applied to SAL models.
7. The particular model needs to be capable of more flexibly learning from experience. In particular we hope to take advantage of the work in ACT-R on combined learning from instruction and observation. With respect to language processing we will need a system that can more robustly respond to instruction and this may require taking advantage of the more continuous and approximate representations that Leabra allows.

This demonstration project has been useful in revealing some challenges that arise in constructing a truly embodied and flexible cognitive architecture. These challenges highlight gaps in existing psychological theory, and suggest that building embodied architectures may be crucial to progress in the field. The project has also established the surprising power and flexibility available with even a simple combination of ACT-R and Leabra architectures. The two architectures are quite compatible despite two different levels of focus, suggesting that these two approaches have converged on a correct overall theory of cognition. The combined SAL architecture is capable of accomplishing tasks that are fundamentally beyond the reach of either ACT-R or Leabra operating in isolation. This success suggests that substantial progress has already been made in understanding the human mind in enough detail to replicate it.

Future Evolution of the SAL Architecture

In addition to the more task-specific challenges facing the SAL model outlined above, there are a number of more general and far-reaching issues that will shape the future evolution of the architecture. Overall, we hope to evolve this architecture from a relationship of mutual codependency between components of two separate systems to a more synthetic combination of the two systems. Below we consider the consequences of such a synthesis for some of the modules.

1. Procedural: The ACT-R production system represents a high functionality system that provides the needed control in SAL. Leabra's basal ganglia model represents a much more detailed system that is closer to the neural realities. In the synthesis we will constrain the production system to reflect that neural reality and in the process actually increase its functionality.

(a) **Action Selection.** In the current ACT-R system, a production will fire only if it matches exactly. Exact matching is not tenable in light of either the biology or the needed functionality. Among the productions that do match in ACT-R, selection of the one to fire is made on the basis of learned reinforcements. This constraint that selection only begins to apply after matching is complete is again not tenable in light of either the biology or the needed functionality. The Leabra system would suggest that a dynamic threshold for matching emerges as a function of a competition among candidate productions – a less than perfect match can fire if it is the best thing currently available. This may be critical for allowing the system to learn new productions by reshaping old ones in new ways, as partial matches based on existing knowledge are co-opted and modified for new tasks.

(b) **Production Learning.** One of the functionally powerful mechanisms in ACT-R is production compilation, by which new productions are created. In a typical example, one production will request a critical piece of information be retrieved from declarative memory, the information will be retrieved, and a second production will act on it. Production compilation replaces this with a single step in which the action is directly taken without retrieval. This is critical to the process by which instructions come to directly control behavior. From the Leabra perspective, this kind of learning involves the development of new representations in both prefrontal cortex and basal ganglia, and with sufficient levels of repetition, may become independent of these systems and be encoded directly between the parietal and motor frontal areas. Thus, there are likely to be important shifts in the locus of neural activity over the course of production learning, which may produce important functional benefits in terms of reducing central capacity bottlenecks for highly practiced procedures. In the context of a navigation and target-search task, the ability to avoid obstacles and perform local navigation may become highly automated and free up more resources for visual search and higher-level route planning.

(c) **Pattern Matching.** One of the critical questions is exactly how complex a pattern can be recognized in a single production cycle or a single pass through the basal ganglia. The simpler production rules in ACT-R can be realized in Leabra-like processes. However, Anderson has identified a more powerful kind of rule involving what he calls dynamic pattern matching which seems critical for human intelligence. In particular, they are critical for learning from instruction and demonstration – the typical means of social communication of knowledge. This can be supported through Leabra's dynamic gating system. The current mechanism in ACT-R is only able to learn dynamic pattern-matching productions from other dynamic-pattern matching productions – it is not able to generalize explicitly matched rules to dynamically matching ones. Considering such mechanisms, in conjunction with Leabra capabilities such as dynamic gating, provides a good area for exploration.

2. Declarative: When ACT-R retrieves a chunk, it selects the most active one. The activation of a chunk reflects its past frequency of occurrence, its strength of association to the current context, and how well it matches the retrieval probe. All of these factors are combined to yield a quantity that reflects the likelihood that the chunk is the desired memory. A series of blending models have been developed in ACT-R for merging the contribution of multiple chunks into a single retrieved memory. In Leabra, there are actually two underlying systems supporting declarative retrieval: the hippocampus and posterior cortex. These systems have different characteristics. The hippocampus behaves more like ACT-R single-chunk retrieval, in that a single coherent chunk is typically retrieved, and it is highly sensitive to context and probe match. However, the posterior cortex can support overlapping distributed representations of multiple chunks at the same time, with each making a graded contribution to the overall memory retrieval process. This is more like the ACT-R blending models. We plan to integrate the Leabra and ACT-R perspectives into a more effective declarative memory.

Part of the effectiveness of declarative memory is the ability to incrementally absorb facts and adjust its generalization threshold to reflect the increasing knowledge base. In Leabra, that property arises from the gradual increase in the size of connection weights as a function of practice. In ACT-R however, while the absolute activation level of chunks increase with practice, their discriminability does not. We have experimented with modifying the ACT-R activation and partial matching equations to more closely reflect the computations in Leabra. This work is an instance of a different sort of integration between ACT-R and Leabra where properties of one are absorbed into the other at a different level of abstraction. This approach is complementary to the integration strategy described earlier and indeed facilitates it.

3. Motor: The current motor actions are issued as discrete requests that are not guided by changing sensory information. A more Leabra-like implementation would have the specific parameters of these actions emerge as a result of a strong constraint-satisfaction process that takes into account many variables (precise location of things, speed of motion, slope of the floor, etc) to produce the desired goal.

4. Visual: The current visual system in ACT-R can only use top-down constraints to select objects to attend to. When these top-down constraints fail to find adequate guidance it is left to select among the objects randomly. There is new work within the ACT-R group on visual salience and how that can provide bottom-up influence. Merging bottom-up and top-down constraints will be critical in the anticipated BICA environments. Leabra provides guidance about how to coordinate the bidirectional top-down and bottom-up effects.

In addition to these module-specific considerations, there are numerous broad-based issues that are common across many different modules. For example, the way that learning is shaped by emotion and motivational states, which in turn are strongly influenced by social interactions, is a critical aspect of human cognition that SAL will need to address more directly. Some of this is captured in the existing reinforcement learning models in each architecture, but these issues really go beyond the confines of the procedural learning system, and shape representations and processing throughout the system.

Conclusion

We are just at the very beginning of what will hopefully be a long and fruitful process of breaking down longstanding barriers between different architectural "camps" in the field, and developing a truly synthetic and powerful understanding of the human cognitive architecture. The joining of forces represented by the SAL team already represents an unprecedented accomplishment of the BICA program, and we look forward to many more. We are confident that the BICA goal of developing a dynamically taskable, adaptive cognitive agent that can be deployed in a wide range of novel environments and task conditions is achievable with the synthesis of ideas represented by our team.

Supplementary Reports on Research Activities

Report 1: Models of Algebra Learning

The ACT-R group has been working on learning in the domain of algebra. They have taken this as a miniature for exploring the taskability issues that will arise in BICA Phase II. An environment has been created for presenting problems and instructions to ACT-R or students much like a common environment is imagined for the BICA Phase II. Empirical explorations have included study of standard algebra taught to children and an isomorph appropriate for instruction to adults. Different studies have looked at learning from typical textbook instruction, learning from examples, and learning from exploration. These reflect the modes of learning that will be required of the BICA agent. Successful ACT-R models have been developed of learning from instruction and learning from examples that are able to predict the learning trajectories of actual students. However, the ACT-R group is still working to characterize the rather remarkable success that students have in learning from discovery. Particularly important for our ongoing efforts to model learning from exploration is understanding how the existing knowledge of students guides their exploration and enables them to interpret the outcomes of this exploration. Again it will be critical in BICA Phase II to be able to characterize the role of prior knowledge in learning about the environment.

One of the important outcomes from this effort was the realization that the current pattern matching in the ACT-R system was not powerful enough to enable processing the abstract relationships in instruction and example. Initial explorations revealed that we could capture the kind of learning students were doing with the more powerful SOAR pattern matcher but that the SOAR pattern matcher was so powerful as to be completely unrealistic biologically – being able to do exponential search in a single match. This led to a restricted version of pattern matching called dynamic pattern matching. One of the early results of interactions between the Leabra and ACT-R research group was the realization that the Leabra dynamic gating system provided a neural model for dynamic pattern matching in ACT-R. One of our future goals is to use this Leabra work to provide a more careful analysis of how dynamic pattern matching should be implemented in ACT-R, what its limitations are, and how to characterize its time costs relative to regular ACT-R pattern matching.

Another aspect of this research has been to look at the learning of algebra in an fMRI scanner. There is a mapping of the ACT-R modules onto specific brain regions and we have been able to use activation in these regions to inform our models of these tasks. As such it represents the potential role of fMRI in BICA Phase II.

The imaging experiment looked at the brain signature obtained while participants performed certain algebraic transformations. It manipulated two factors. One was whether the transformations were relatively simple algebraically (e.g., 3*4X=24 --> 12X=24) or relatively more complex (e.g., 3*(4 + x)=24 --> 12+3x=24). The second was whether this was early or late in the learning. Figure 6 illustrates the results obtained from four cortical regions. The dotted lines connect the actual data and the solid lines are the predictions of the theory.

(a) Figure 6a shows the response in the region of the motor cortex that controls hand movement and corresponds to the manual modules. Since it required more hand motions to execute complex transformations there is greater activation in this region for complex transformations. However, the number of hand movements did not vary early to late. While participants took longer earlier and so the response is stretched over a greater time span, the total area under the curves is the same.
(b) Figure 6b shows the response in the region of the fusiform gyrus, which we have found to tap the high-level activity of the visual module which is all ACT-R represents. This shows both greater activation for complex equations and greater activation early. Like the motor region the greater activation for more complex equations is predicted because more encoding is required to enable a complex transformation. The ACT-R model does not really predict the learning effect. The solid lines in Figure 6b were only produced by the ad hoc reduction of encoding time for the late curves. This points to a place were a SAL model with Leabra-based visual learning would do better.
(c) Figure 6c represents activity in a prefrontal region that many researchers have found to

reflect retrieval from declarative memory, presumably because it holds controlling retrieval cues. The pattern we see here is one that has been observed in almost every study we have done manipulating complexity and practice – which is large effects of both. An effect of complexity is predicted because more retrievals are required for more complex equations and an effect of practice is predicted because the major dimension of learning in the ACT-R model is the drop out of some declarative retrievals and a reduction in the time to perform others.

(d) Figure 6d shows the activation pattern in the anterior cingulate cortex, which is believed to reflect control operations. Like the other regions it shows greater activation for the more complex condition but the effect of practice is complicated. In the case of simple equations there is no effect on area under the curve but in the case of complex equations there is actually an increase in area under the curve with practice late in the performance of the transformation. This reflects the fact that mastering later transformations makes students sensitive to decisions about signs that they ignored earlier. These sign decisions come late in entering the transformation.

These data illustrate how imaging data can confirm the theory in ways that range from expected to surprising (e.g., Figure 6a to Figure 6c to Figure 6d) and at the same time indicate places where the analysis needs to be modified (i.e., Figure 6b).

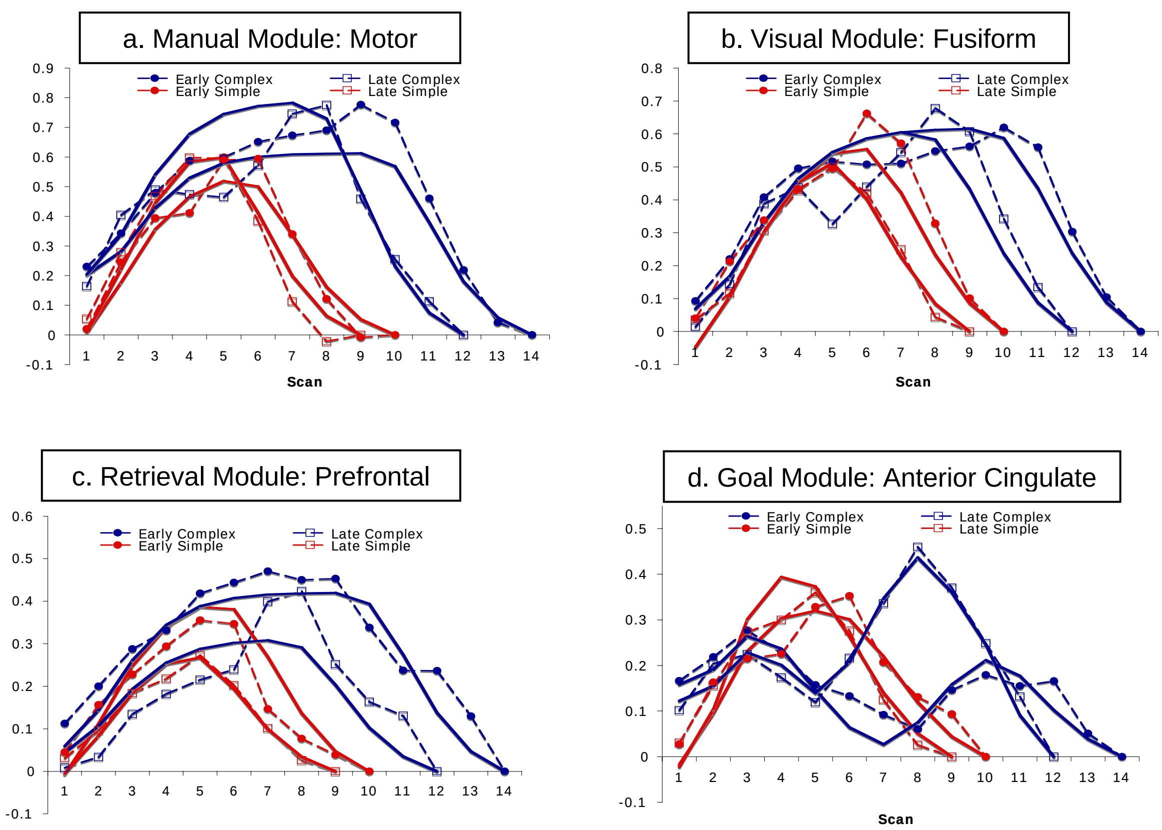
a. Manual Module: Motor
Early Complex
Early Simple
Late Complex
Late Simple
Scan
b. Visual Module: Fusiform
Early Complex
Early Simple
Late Complex
Late Simple
Scan
c. Retrieval Module: Prefrontal
Early Complex
Early Simple
Late Complex
Late Simple
Scan
d. Goal Module: Anterior Cingulate
Early Complex
Early Simple
Late Complex
Late Simple
Scan

Figure 6. BOLD response in four cortical regions for simple and complex equations, early and late in practice. Dotted lines connect actual data and solid lines are predictions from ACT-R modules.

Report 2: Attentional Blink Model

To demonstrate that the SAL architecture can model behavioral phenomena that are difficult or impossible to model with either ACT-R or Leabra alone, we developed a SAL model of the Attentional Blink phenomenon.

In the experiment associated with Attentional Blink, participants are presented with rapid streams of 20 characters (at a presentation rate of 100 ms/character), most of which are digits (distracters), and some (0, 1 or 2) are letters (targets). The goal for the participant is to identify and report the targets (the letters). In streams with two targets, response accuracy differs depending on how many distracters are presented between the two targets. This distance is called the lag, where lag 1 means no distracters in between the targets, lag 2 one distracter, etc. If the lag is five or more, accuracy on both targets is the same, around 80%. When the lag is less, accuracy on the second target is worse then on the first target with one exception: when the lag is one, meaning the targets are immediately sequential, accuracy on both targets is again identical. However, in that case participants make a different error: they sometimes report the two targets in the wrong order.

To model the task we used a SAL prototype that is architecturally similar to Figure 3, where the “Vision what” module of ACT-R was replaced by a modified Leabra vision model. This modified Leabra model does not reset its activations between stimuli, and its output is a set of graded activation values at each time step, rather than a final symbolic determination. Having been previously trained on the character set used in the human experiment, the Leabra network was presented with the stream of characters that the participants also perceived. Due to the speed of the input, the network was not always able to reach peak activity for a particular classification, and would sometimes be in a transition in between two classifications when queried by the ACT-R part of the model. Figure 7 shows a sample graph of output activation of the vision module.

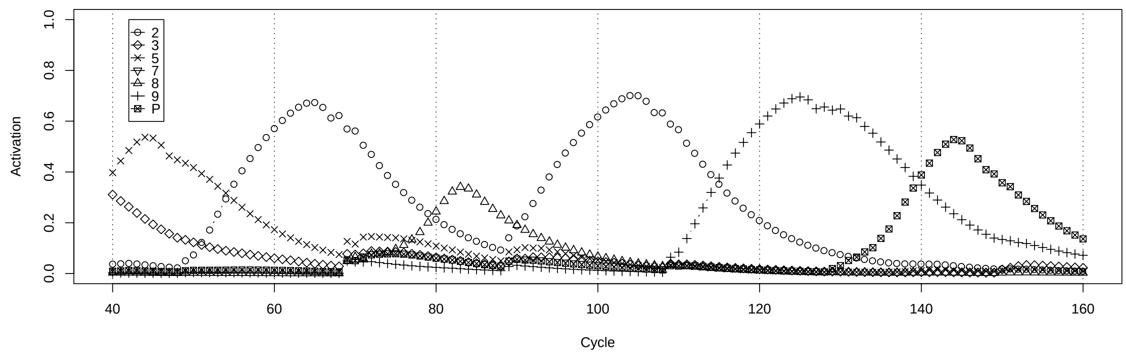

Figure 7. Example of activation rise and fall for 7 output cells (for characters 2, 3, 5, 7, 8, 9 and P). Vertical lines indicate where a new character was presented to the network. Note that it takes around 10 cycles (20 cycles maps onto 100 ms) before a new character produces a rise in activation in the output layer. If the network is sampled at a moment when multiple characters are active (e.g., 2 and 9 at cycle 115) it will pass all candidates on to ACT-R.

The Leabra component not only provides a realistic account of how a letter is processed, but also explains why the order of targets is sometimes mixed up when the two targets are right after each other: if the network perceives two characters at the same time it has no way to deduce the order, so has to make a random guess (participants report that it seems like the two letters are superimposed on each other).

Once Leabra has recognized characters, ACT-R has to classify and possibly memorize them. The ACT-R model assumes that targets are stored in the imaginal buffer. In order to determine whether something is a target, its category (letter or digit) has to be retrieved from declarative memory first. The rapid presentation rate puts a heavy strain on the capacity of the architecture to keep track of all the input. Once a first target has been detected, the additional task of storing it creates additional load on the system, creating a ripple effect (similar to a traffic jam) on processing further down the line. It turns out that the peak of this "cognitive traffic jam" is 200 to 300 ms after the presentation of the first target, which is exactly where the attentional blink effect peaks. Figure 8 shows the comparison between model and data.

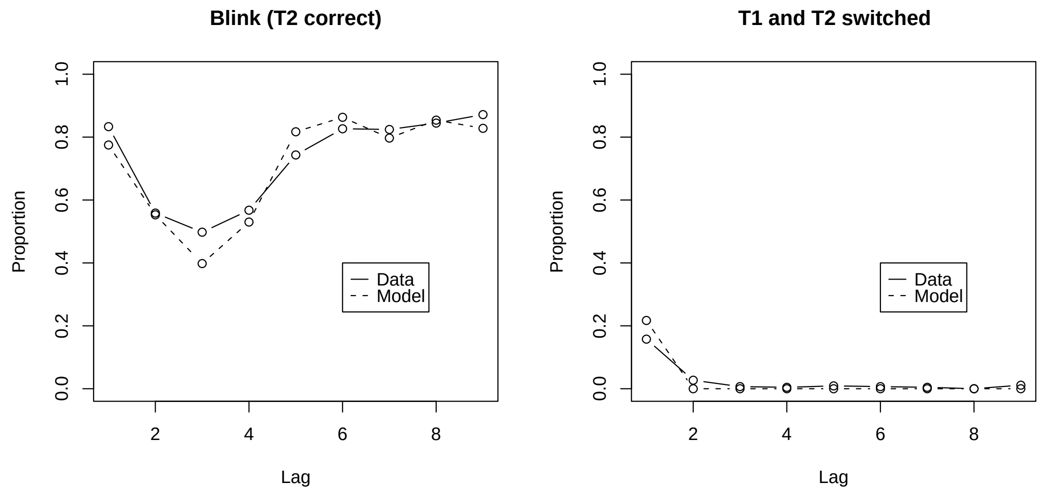

Figure 8. Data and model results of the attentional blink experiment. Left: proportion of correct T2 responses, showing the attentional blink effect. Right: proportion of switch errors, where both targets are named correctly but in reverse order.

The model of attentional blink combines the strengths of the two architectures: the fine-grained perceptual capabilities of the Leabra architecture can explain the reversals at lag 1 (Figure 8, right), and ACT-R's serial processing constraints on individual modules can explain the blink effect (Figure 8, left).

Report 3: Additional ACT-R Research Thrusts

While for phase I the ACT-R part of this effort has been concentrated at CMU, it also included a number of consultants who have been longstanding members of the ACT-R community and are envisioned to play a much more substantial role in Phase II to help us tackle the difficulty and complexity of the tasks ahead. Those consultants involve 12 people distributed over 8 institutions:

Air Force Research Laboratory Mesa
Naval Research Laboratory Washington DC
Rensselear Polytechnic Institute
Drexel University
Rice University
Pennsylvania State University
Xerox Palo Alto Research Center
University of Michigan

In addition, Alion Science & Technology has also been involved as a subcontractor.

During phase I, those consultants have investigated architectural issues of central interest to the goal of BICA and in particular to the challenge tasks and environments that have been discussed. To avoid scattering our efforts and to bring significant resources to bear on those problems, they have organized into larger teams focused on a small number of research questions that are both of fundamental scientific interest and directly relevant to the BICA challenge tasks. The teams considered constraints ranging from functional performance to neuroscience evidence in evaluating architectural designs for the hybrid SAL architecture. Those research questions, together with the teams focused on them, are:

Spatial modules and navigation (AFRL, NRL, Rice)
This research thread focuses on developing a new spatial module for SAL. The current ACT-R visual module only provides a relatively flat and static view of the world and significant enhancements are required to operate effectively in the 3-dimensional dynamic world of the challenge problems as well as to reflect a more complete picture of the neuropsychological evidence regarding human spatial abilities. This new module would work in concert with the improved and integrated versions of the visual and motor modules and would provide competencies including support for frames of reference, mental imagery, magnitude estimations, spatial transformations and navigation.

Situation Awareness/Multimodal Integration/Episodic Memory (Drexel, RPI, AFRL)
This research thread focuses on developing memory competencies that integrate external experiences across time and across sensory modalities. The current ACT-R declarative memory represents information of different points in time or sensory modalities as independent chunks without any links to a common integrated picture of the environment associated with episodic memory capacities. This research effort has explored ways of maintaining robust continuous situation awareness in a dynamic environment by developing memory representations and processes that integrate experience across time and sensory modalities. These do not take the form of a new architectural module but instead consist of an elaboration of the existing memory representation and retrieval processes.

Motor modules/Robotic Embodiments (Rice, PSU, NRL)
This research thread focuses on developing motor modules consistent with the simulated embodiment of the BICA agent, including simplified lower-body, articulated upper body and complex manipulators. The current ACT-R motor module is limited to two-handed keyboard actions and needs considerable generalization to support the proposed embodiment. This research thread also investigates and develops architectural assumptions such as direct perceptual-motor module links and continuous control of motor actions. Because of the importance on learning and development and the difficulty in programming controllers for complex activators, this thread also focuses on learning mechanisms for the new motor module in concert with cognitive learning mechanisms such as production compilation. This coordinated learning across architectural modules constitutes a significant challenge and a major innovation.

Language/Ontologies (Michican, Xerox PARC, Alion)
This research thread focuses on developing language capabilities to support the dialogue with teacher(s) and fellow BICA agents in the proposed task environment. These include support for understanding instructions, language acquisition similar to human development, and support for the role of language in cognition. This thread also focuses on representational issues, such as the adoption of common ontologies for integrating models.

Social Interactions/Theory of Mind (Alion, Drexel)

This research thread focuses on the social capabilities involved in interacting with teacher(s) and BICA agents in the task environment, including inferring and understanding the beliefs, intentions and actions of others, sharing mental pictures, and working in cooperation and competition. This thread also focuses on understanding the role of mirror neurons in learning from demonstration and imitation. This will require integration both with the visual/spatial modules to provide a representation comparable to that provided by mirror neurons and with the motor modules to allow the closing of the loop from visual input to motor action and back. This could take the form of a separate intentional module, of direct connections between perceptual and motor modules, of cognitive skills modulating those interactions, or any combination of the above

All of these research directions involve interaction between with the teams (e.g. the second topic will need to take as input the output of the spatial modules developed under the first topic) as well as with the CMU team, requiring continuous attention to integrating these efforts conceptually (as well as in software) into a coherent framework.

Report 4: Additional Leabra Research Thrusts

In addition to its contribution to the SAL effort discussed previously in this report, the Leabra group also developed Leabra-specific architectural elaborations during Phase I, including new proposed models of brain regions that we have not previously addressed, as well as the adaptation and “scaling up” of existing models. Although these elaborations were developed as part of our independent Phase I effort, we anticipate that many of them will fit naturally into the SAL architecture. Consistent with past approach, these architectural elaborations are based closely on the intersection of (a) what is known about how the brain solves these same problems, and (b) the principles of successful computational models, derived from prior Leabra research as well as through the literature. Some of these key principles are reviewed at the end of this report. Figure 9 provides an illustrative overview of these architectural elaborations, including both existing and planned network models, in one possible configuration in the context of the proposed SAL architecture.

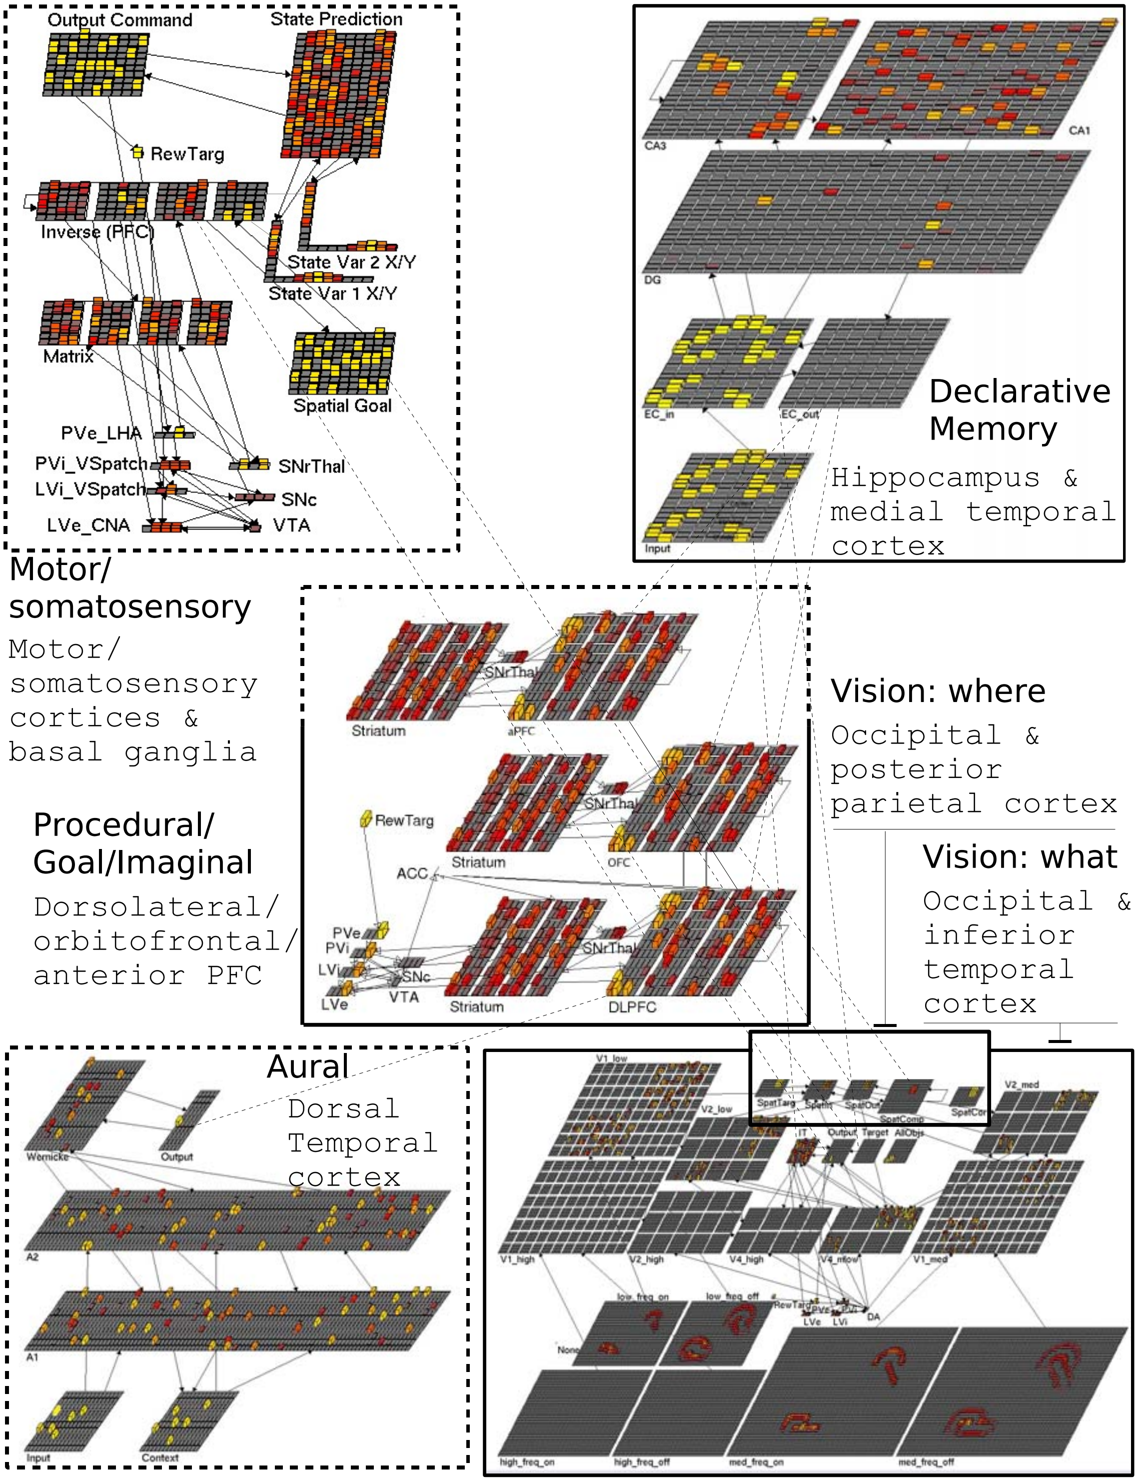

Figure 9. Planned Leabra networks for the SAL architecture. See text for details. Models with solid borders are planned elaborations of existing models; dashed borders indicate models which have been planned in phase I, but do not yet exist in working form. Dotted lines denote planned connections between models.

Module 1: Visual "What" – Ventral Visual Stream

To date, the bulk of our effort to apply Leabra to real-world problems has been focused on models of the visual system. Our current ventral stream model performs well on object identification and classification compared to state-of-the-art statistical learning methods, despite a relatively brief development and training period. For example, it can identify 100 different 3D objects at up to 3x scale variations and 36 degree rotations, at 95% accuracy, and exhibits excellent performance on real-world photographs. In addition, models developed initially for object recognition show emergent visual search behavior; see the "Visual Where" section below. The success of these visual system models indicates that rapid progress could be made in constructing large-scale models of other brain systems; see the final section of this report, "the Leabra approach"

Our existing ventral stream model is a Leabra network based on the known structure of the visual system. It has a hierarchical set of layers representing areas V1, V2, V4 and IT. One key biologically-inspired principle in this and all Leabra models is the inclusion of abundant feedback connections. These connections allow for biologically realistic error driven learning, and allows a type of constraint satisfaction that greatly enhances performance over feed-forward models. A second key principle of this model is mimicking of the collapsing receptive field structure of the biological visual system. Each unit is constrained to receive from a limited subset of those in the previous layers, so that the size of spatial receptive fields increases in progressively higher layers. After training with a combination of associative and error-driven learning, units' response properties resemble those found in the mammalian system: higher units respond to increasingly complex object features, and with increasing spatial invariance. These principles are commonly thought to play similarly important roles in other sensory systems.

We anticipate that adding input filters for color and texture, both of which serve as powerful cues to object location and identity, will improve performance further. We have also architected and begun work on a simplified semantic network to represent parts of objects as well as objects. This function will work synergistically with ACT-R's sophisticated control process. Finally, improvements in the dorsal stream components described next will allow the network to better focus on an object or feature of interest and ignore background clutter.

Module 2: Visual "Where" – Dorsal Visual Stream

In the Leabra model of the visual system, the dorsal and ventral streams are both integrated into a single interactive network. The network layers that perform spatial location and attention leverage the limited spatial information from the ventral stream ("what") layers. The same reciprocal connections that enable biologically realistic error-driven learning in the ventral stream also allow top-down object-based attention: activating the output representation of a target causes the lower-level units representing features (and approximate location) of that object to gain an advantage in the ongoing competition with surrounding units. Competition within the spatial layers then produces enhanced activity at the location of the target object. (This process, known as biased competition, is a prominent explanation of attention that can explain a large variety of sensory and cognitive attentional phenomena). Because of inhibitory competition, this attention on important features also *reduces* the activation of irrelevant aspects of stimuli, improving both visual search and object recognition performance.

This model is able to localize targets in visual fields with multiple objects with up to 95% accuracy for simple objects, and 76% for more complex objects. This effect of object complexity on visual search ability is a well-documented feature of the human visual system. Ongoing work on this model includes adding motion signals (also processed by the dorsal stream) and integrating it with other sensory modalities (e.g., auditory and somatosensory location signals).

Module 3: Auditory Processing – Dorsal Temporal Lobe
During Phase I we designed an architecture for audition based on the strategy used in the visual system: apply the Leabra algorithm in a network modeled after the known architecture of human auditory brain areas. Our approach to auditory processing is to filter the input over time as well as frequency dimensions, consistent with what is known to occur via the pathway from hair cells through sensory thalamus. Input neurons will code for instantaneous frequencies and for particular changes in frequencies over time. A second input layer provides the same information for the previous time step, allowing context information about the previous sound to guide processing, as in human speech recognition. Thereafter, the same hierarchical approach as employed in visual processing will be applied: each unit will receive from lower units representing a pre-specified range of frequency and frequency changes, and the same learning algorithm will be applied so that they form more specific receptive properties according to their usefulness in identifying words and sounds.

Modules 4, 5, 6: Procedural, Goal, Imaginal - Prefrontal Cortex/Basal Ganglia (PFC/BG)
The Leabra architecture for the prefrontal cortex and basal ganglia evolved in significant ways during Phase I. This system is specialized for working memory and executive control, through its ability to actively maintain (PFC) and adaptively update (BG) task-relevant information (goals, partial products of ongoing processing, etc). As noted earlier, the gating provided by the BG is functionally similar to the procedural module in ACT-R, in that it activates specific memory or control representations in PFC when certain conditions are met. These PFC representations play the role of the goal and imaginal buffers in ACT-R, providing information for further procedural decisions. In addition, the maintained representations exert a direct top-down biasing influence to direct attention in sensory systems, and direct the motor system toward specific goals.

Based on known anatomical and electrophysiological characteristics of the brain, the model is organized into micro-anatomical "stripes," each capable of maintaining isolated pieces of information. Each stripe is also *selectively* updateable by mechanisms driven by the basal ganglia. Critically, cells of the basal ganglia actually *learn* to control this selective gating of active maintenance in the PFC under the influence of reward prediction error signals delivered by ventral tegmental area (VTA) and substantia nigra pars compacta (SNc) dopamine cells. These dopamine cells, in turn, are themselves driven by a distinct subsystem (Primary Value, Learned Value, or PVLV) involving learning in the amygdala and ventral striatum. In this system, global working memory representations can be adaptively modified over time to exert flexible control.

During Phase I, we elaborated several improvements to the PFC/BG system. These (shown in Figure 9) include: (1) sub-specialization among the various areas of the PFC, e.g., the anterior pole (aPFC) for higher abstract task switching and the olfactory cortex (OFC) for goal and value representations; (2) integration of related areas known to play a role in cognitive control, e.g., the anterior cingulate cortex (ACC) for conflict monitoring and error expectation, which regulates the maintenance and updating of PFC representations; (3) output gating (in addition to existing maintenance gating) to control when maintained information influences processing and motor output; and (4) integration of the PFC/BG working memory system with a similar BG-based functionality known to exert control over motor behaviors, especially involving motor plans.

Module 7: Declarative Memory - Hippocampus/Medial Temporal Lobe (HC/MTL)
As noted earlier, the hippocampus and medial temporal lobe function as a sub-system specialized for the rapid learning of arbitrary information. The hippocampus receives information from a wide range of cortical areas, essentially representing the entire cortical state at one time in the entorhinal cortex (EC) input layer of the hippocampus. It encodes this state using very sparse representations (in the dentate gyrus (DG) and area CA3), which produce pattern separation to avoid interfering with other representations. This hippocampal state can be recalled from a partial cue (via recurrent collaterals in CA3), and this recall spreads back from the hippocampus (via invertible CA1 representations connected to EC) out to the cortex, resulting in the reinstatement of the original cortical representation. This general ability is critical for episodic memory and navigation.

Our hippocampal model is well established and tested on a wide range of tasks. The next step is to integrate it with our visual model, to enable rapid learning of specific spatial locations and object-name associations.

Module 8: Motor/Somatosensory
We developed some initial models of motor control and somatosensory feedback. The initial model learns to control a two joint arm, but the principles should generalize to more complex effectors, because everything is based on learning. Our goal is that an untrained model following this architecture could be placed in a body with different sensors and effectors, and would learn to manipulate its motor systems.

The inputs to the motor model are a goal location and somatosensory information on the relevant system's current state, such as position and joint angles. The "inverse" system then computes a motor command to move toward the goal from the current position. This system learns through rewards (successful reaching to the goal or not). This learning signal is aided by a predictive layer that anticipates the sensory result of a given motor command (i.e., a "forward" model). This predictive system learns more quickly since it has the benefit of an error signal at every time step. Our next step is to integrate this model with the PFC/BG system to provide better action selection abilities (especially for temporal sequences of actions), and with the visual-spatial model to provide better perceptual representations of motor effector state.

Software Development
We are nearly done with a major overhaul of the PDP++ neural simulation software that we have developed and used since 1993. This new version provides much greater flexibility in the way

that Leabra models can be trained and tested, by adopting a graphical programming language approach similar to that of widely-used experimental testing software (e.g., E-Prime). It also provides a much more powerful and flexible GUI, with tabbed browser functionality, and 3D visualization based on the OpenGL API. These changes will enable models to interact more realistically and continuously with a virtual or real world, as well as interchanging information with ACT-R under the SAL architecture, without the need of modelers to change the core code of Leabra.

Background info on the Leabra Framework

All of the above models use a common set of computational mechanisms that reflect a long-term effort to integrate biological, computational, and cognitive constraints on the fundamental nature of neural processing in the cortex and other brain areas. Perhaps the most important mechanisms are the combined Hebbian and biologically-plausible error-driven learning mechanisms, which enable the models to self-organize new representations to solve challenging tasks. We have repeatedly found that using both forms of learning, which is unique to the Leabra framework, is critical for many learning domains. These learning mechanisms are substantially more effective when combined with inhibitory competition, which drives different neurons to specialize for representing different information. This inhibitory dynamic, captured in Leabra using a computationally-efficient k-winners-take-all (kWTA) mechanism, is also critical for stabilizing the activation dynamics that emerge with bidirectionally-connected networks. This bidirectional connectivity is critical for biologically-plausible error-driven learning, and for capturing interactions between bottom-up sensory-driven processing and top-down goal-driven processing (as in our visual search model). Most other neural networks do not support full bidirectional processing because it produces complex activation dynamics that are difficult to control without a solid inhibitory mechanism like the kWTA present in Leabra. Finally, Leabra relies extensively on coarse-coded distributed representations, which enable generalization to novel inputs and are highly efficient and robust. All of these features together produce a synergistic emergent dynamic that seems critical for capturing the magic of human cognition.